AD-A171 032 NAVAL C3 DISTRIBUTED TACTICAL DECISION MAKING(U) 1/1
MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR
INFORMATION AND DECISION SYSTEMS M ATHANS ET AL.
UNCLASSIFIED 27 JUL 86 LIDS-IR-1582 N00014-84-K-0519 F/G 5/10 NL

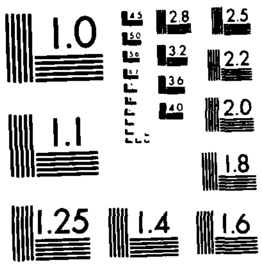
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0

AD-A171 032

DTIC FILE COPY

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) NAVAL C^3 DISTRIBUTED TACTICAL DECISION MAKING		5. TYPE OF REPORT & PERIOD COVERED 1 April 1986 - 30 June 1986
		6. PERFORMING ORG. REPORT NUMBER LIDS-IR-1582
7. AUTHOR(s) Michael Athans, Alexander H. Levis		8. CONTRACT OR GRANT NUMBER(s) N00014-84-K-0519
9. PERFORMING ORGANIZATION NAME AND ADDRESS Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 649-003
11. CONTROLLING OFFICE NAME AND ADDRESS ENGINEERING PSYCHOLOGY PROGRAMS (CODE 442EP), Office of Naval Research, Arlington, Virginia 22217		12. REPORT DATE July 27, 1986
		13. NUMBER OF PAGES i+25
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) Mr. J. R. Simpson (Code 111SP), Mathematical Sciences Division, Office of Naval Research, Arlington, Virginia 22217		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release, distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

AUG 8 1986

A

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Command and Control, Distributed Decisionmaking, Organization Theory, Petri Nets

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Progress on eight research problems addressing distributed tactical decisionmaking is described.

LABORATORY FOR INFORMATION AND DECISION SYSTEMS

Massachusetts Institute of Technology
Cambridge, Mass., 02139, U.S.A.

OSP NUMBER 95442 LIDS-IR-1582

QUARTERLY REPORT

for the period

1 April 1986 to 30 June 1986

for

NAVAL C^3 DISTRIBUTED TACTICAL DECISION MAKING

Contract Number N00014-84-K-0579
Work Unit NR 649-003

Submitted to:

Mr. J. R. Simpson, Code 111SP (3 copies)
Dr. R. Madan, Code 414
Dr. W. R. Vaughan, Jr., Code 442 EP

Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Submitted by:

Michael Athans
Alexander H. Levis

July 25, 1986

1

SERVOMECHANISMS LABORATORY
1939-1959
ELECTRONIC SYSTEMS LABORATORY
1959-1976

NAVAL C^3 DISTRIBUTED TACTICAL DECISIONMAKING

1. PROJECT OBJECTIVES

The objective of the research is to address analytical and computational issues that arise in the modeling, analysis and design of distributed tactical decisionmaking. The research plan has been organized into two highly interrelated research areas:

(a) Distributed Tactical Decision Processes; and
(b) Distributed Organization Design.

The focus of the first area is the development of methodologies, models, theories and algorithms directed toward the derivation of superior tactical decision, coordination, and communication strategies of distributed agents in fixed organizational structures. The framework for this research is normative.

The focus of the second area is the development of a quantitative methodology for the evaluation and comparison of alternative organizational structures or architectures. The organizations considered consist of human decisionmakers with bounded rationality who are supported by C^3 systems. The organizations function in a hostile environment where the tempo of operations is fast; consequently, the organizations must be able to respond to events in a timely manner. The framework for this research is descriptive.

2. STATEMENT OF WORK

The research program has been organized into seven technical tasks - four that address primarily the theme of distributed tactical decision processes and three that address the design of distributed organizations. An eighth task addresses the integration of the results. They are:

2.1 Real Time Situation Assessment

Static hypothesis testing, the effect of human constraints and the impact of asynchronous processing on situtation assessment tasks will be explored.

2.2 Real Time Resource Allocation

Specific research topics include the use of algebraic structures for distributed decision problems, aggregate solution techniques and coordination.

2.3 Impact of Informational Discrepancy

The effect on distributed decisionmaking of different tactical information being available to different decisionmakers will be explored. The development of an agent model, the modeling of disagreement, and the formulation of coordination strategies to minimize disagreement are specific research issues within this task.

2.4 Constrained Distributed Problem Solving

The agent model will be extended to reflect human decisionmaking limitations such as specialization, limited decision authority, and limited local computational resources. Goal decomposition models will be introduced the derive local agent optimization criteria. This research will be focused on the formulation of optimization problems and their solution.

2.5 Evaluation of Alternative Organizational Architectures

This task will address analytical and computational issues that arise in the construction of the generalized performance-workload locus. This locus is used to describe the performance characteristics of a decisionmaking organization and the workload of individual decisionmakers.

2.6 Asynchronous Protocols

The use of asynchronous protocols in improving the timeliness of the organization's response is the main objective of this task. The tradeoff between timeliness and other performance measures will be investigated.

2.7 Information Support Structures

In this task, the effect of the C^3 system on organizational performance and on the decisionmaker's workload will be studied.

2.8 Integration of Results

A final, eighth task, is included in which the various analytical and computational results will be interpreted in the context of organizational bounded rationality.

3. STATUS REPORT

In the context of the first seven tasks outlined in Section 2, a number of specific research problems have been formulated and are being addressed by graduate research assistants under the supervision of project faculty and staff. Research problems which were completed prior to or were not active during this last year have not been included in the report.

3.1 DISTRIBUTED TEAM HYPOTHESIS TESTING WITH SELECTIVE COMMUNICATIONS

Introduction

A major research objective has been completed and documented in the MS thesis of Jason D. Papastavrou under the supervision of Professor Michael Athans. In this research we formulated, solved, and analyzed a distributed hypothesis testing problem which is an abstraction of a wide class of distributed team decision problems. It represents a normative version of the "second-opinion" problem in which a primary decision maker (DM) has the option of soliciting, at a cost, the opinion of a consulting DM when faced

with an ambiguous interpretation of uncertain evidence.

<u>Motivating Examples</u>

Our major motivation for this research is provided by generic hypothesis testing problems in the field of Command and Control. To be specific, consider the problem of target detection formalized as a binary hypothesis testing problem (H_0 means no target, while H_1 denotes the presence of a target). Suppose that independent noisy measurements are obtained by two geographically distributed sensors (Figure 1). One sensor, the <u>primary</u> DM, has final responsibility for declaring the presence or absence of a target, with different costs associated with the probability of false alarm versus the probability of missed detection. If the primary DM relied only on the measurements of his own sensor, then we have a classical centralized detection problem that has been extensively analyzed. If the actual measurements of the second sensor were communicated to the primary DM, we have once more a classical centralized detection problem in which we have two independent measurements on the same hypothesis; in this case, we require communication of raw data and this is expensive both from a channel bandwidth point of view and, perhaps more importantly, because radio or acoustic communication can be intercepted by the enemy.

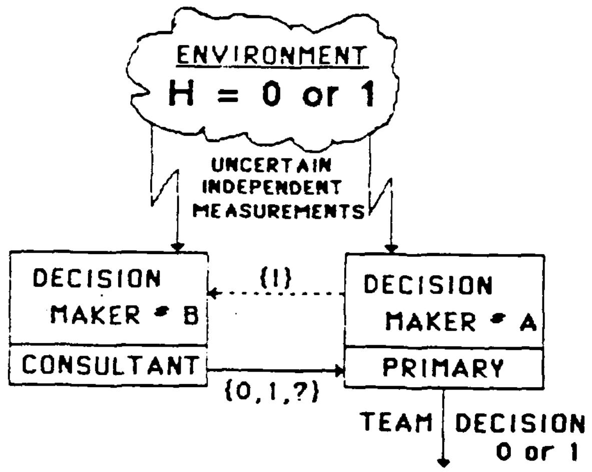

Figure 1. Problem Formulation

Continuing with the target detection problem, we can arrive at the model that we used in the research by making the following assumptions which model the desire to communicate as little as possible. The primary DM can look at the data from his own sensor and attempt to arrive at a decision using a likelihood-ratio test (ltr). Quite often the primary DM can be confident about the quality of his decision. However, we can imagine that there will be instances that the data will be close to the decision threshold, corresponding to an ambiguous situation for the primary DM. In such cases it may pay off to incur a communications cost and seek some information from the other available sensor. It is important to establish what is the nature of the information to be transmitted back to the primary DM.

In our research, we assume the existence of a consulting DM having access to the data from the other sensor. We assume that the consulting DM has the ability to map the raw data from his own sensor into decisions. The consulting DM is "activated" only at the request of the primary DM. It is natural to speculate that his advise will be ternary in nature: YES, I think there is a target; NO, I don't think there is a target; and SORRY, NOT SURE MYSELF. Note that these transmitted decisions in general require less bits that the raw sensor data, hence the communication is cheap and more likely to escape enemy interception. Then, the primary DM, based upon the message received from the consulting DM, has the responsibility of making the final binary team decision on whether the target is present or absent.

The need for communicating with small-bit messages can be appreciated if we think of detecting an enemy submarine using passive sonar (Figure 2). We associate the primary DM with an attack submarine, and the consulting DM with a surface destroyer. Suppose that both have towed-array sonar capable of long-range enemy submarine detection. Request for information from the submarine to the destroyer can be initiated by having the submarine sonar emit a low power sonar pulse. A short coded sonar pulse can be used to transmit the recommendation from the destroyer to the submarine. Thus, the submarine has the choice of obtaining a "second opinion" with minimal compromise of its covert mission.

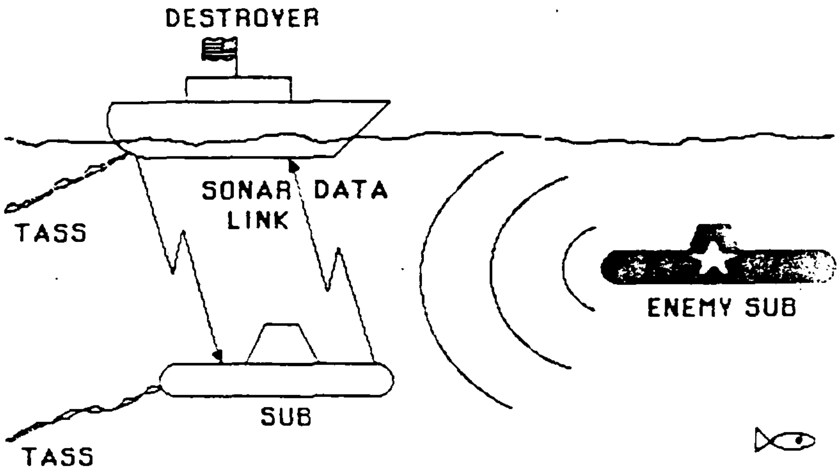

Figure 2. Anti-Submarine Warfare (ASW) Example

Of course, target detection is only an example of more general binary hypothesis-testing problems. Hence, one can readily estend the basic distributed team decision problem setup to other situations. For example, in the area of medical diagnosis we imagine a primary physician interpreting the outcomes of several tests. In case of doubt, he sends the patient to another consulting physician for other tests (at a dollar cost), and seeks his recommendation. However, the primary physician has the final diagnostic responsibility. Similar scenarios occur in th intelligence field where the "compartmentalization" of sensitive data, or the protection of a spy, dictate infrequent and low-bit communications. In more general military Comamnd and Control problems, we seek insight on formalizing the need to break EMCON, and at what cost, to resolve tactical situation assessment ambiguities.

Prior Research

The solution of distributed decision problems is quite a bit different, and much more difficult, as compared to their centralized counterparts. Indeed there is only a handful of papers that deal with solutions to distributed hypothesis-testing problems. The first attempt to illustrate the

difficulties of dealing with distributed hypothesis-testing problems was published by Tenney and Sandell; they point out that the decision thresholds are in general coupled. Ekchian and Tenney deal with detection networks in which downstream DM's make decisions based upon their local measurements and upstream DM decisions. Kushner and Pacut introduced a delay cost (somewhat similar to the communications cost in our model) in the case that the observations have exponential distributions, and performed a simulation study. Recently, Chair and Varshney have pointed out how the results by Tenney and Sandell can be extended in more general settings. Boettcher and Tenney have shown how to modify the normative solutions to reflect human limitation constraints, and arrive in at normative/descriptive model that captures the constraints of human implementation in the presence of decision deadlines and increasing human workload; experiments using human subjects showed close agreement with the predictions of their normative/descriptive model. Finally Tsitsiklis and Athans demonstrate that such distributed hypothesis-testing problems are NP-complete; their research provides theoretical evidence regarding the inherent complexity of solving optimal distributed decision problems as compared to their centralized counterparts (which are trivially solvable).

Contributions of this Research

The main contribution of this research relates to the formulation, quantification, and optimal solution of the team decision problem described above. Under the assumption that the measurements are conditionally independent, we show that the optimal decision rules for both the primary and the consulting DM are deterministic and are expressed as likelihood-ratio tests with constant thresholds which are tightly coupled.

When we specialize the general results to the case that the observations are linear and the statistics are Guassian, then we are able to derive explicit equations for the decision thresholds for both the primary and consulting DM. These thresholds equations are very nonlinear and tightly coupled, thereby necessitating an iterative solution. They provide clear-cut

evidence that the DM's indeed operate as team members; their optimal thresholds are very different from those that they would use in isolation, i.e., in a non-team setting.

Numerical sensitivity results for the linear-Guassian case provide much needed intuitive understanding of the problem and concrete evidence that the team members operate in a more-or-less intuitive manner, especially after the fact. We study the impact of changing the communications cost and the measurement accuracy of each DM upon the decision thresholds and the overall team performance. In this manner we can obtain valuable insight on the optimal communication cost increases, the frequency of communication (and asking for a second opinion) descreases, and the team performance approaches that of the primary DM operating in isolation. In addition, we compare the overall distributed team performance to the centralized version of the problem in which the primary DM had access, at no cost, to both sets of observations. In this manner, we can study the degree of inherent performance degradation to be expected as a consequence of enforcing the distributed decision architecture in the overall decision making process.

Deterioration of the quality of the observations of the primary decision maker result in two very different situations. If the cost of declaring "defacto" the a priori most probable hypothesis is less than the sum of the communication cost plus the cost of the consulting decision maker making the final decision, then the primary decision maker declares the more probable hypothesis to be true. Otherwise, the primary decision maker decides to incur the communication cost and passes the responsibility of the team decision to the consulting decision maker.

When the quality of the observations of the consulting decision maker decreases, less information is requested by the primary decision maker. moreover, the consulting becomes more willing to make the final decision, as he realizes that the primary must be really confused, since the primary is willing to incur the communication cost for information of lesser quality.

The effects of increasing communication cost are very similar to the effects of decreasing quality of the observations of the consulting decision maker, since in both cases the consultant's information becomes less helpful for the team.

Finally, we studied the team performance degradation when one of the team members, either the primary or the consulting DM, has an erroneous estimate of the hypotheses prior probabilities. This corresponds to mildly different mental models of the prior situation assessment. As expected the team performance is much more sensitive to misperceptions by the primary DM as compared to similar misperceptions by the consulting DM. This implies that, if team training reduces misperceptions on the part of the DM's, the greatest payoff is obtained in training the primary DM.

In closing we would like to emphasize the two most important conclusions of this research. First, the optimal decision rules of the two decision makers are coupled. That is the optimal decision rule of each team member depends on the decision rules of the rest of the members of the team. Second, because of this, a team member can make decisions which are in total contrast with the decisions that the same DM would make, if he were to make the final decision alone and not as the part of a team.

Documentation: J. D. Papastavrou, "Distributed Detection with Selective Communications" MS Thesis, Dept. of EECS, MIT, May 1986; also Report LIDS-TH-1563, MIT, May 1986.

J. D. Papastavrou and M. Athans, "A Distributed Hypotheses-Testing Team Decision Problem with Communications Cost" LIDS-P-1538, MIT, February (this paper was presented at the 9th MIT/ONR Workshop on C^3 Systems, Monterey, CA, June 1986). Also submitted to 25th IEEE Conference on Decision and Control.

Presentations: Partial results were presented at the 8th MIT/ONR C^3 Workshop, June 1985. The complete results were presented at the 9th C^3 Workshop, June 1986.

3.2 DISTRIBUTED HYPOTHESIS TESTING WITH MANY AGENTS

Background: The goal of this research project is to develop a better understanding of the nature of the optimal messages to be transmitted to a central command station (or fusion center) by a set of agents who receive different information on their environment. In particular, we are interested in solutions of this problem which are tractable from the computational point of view. Progress in this direction has been made by studying the case of a large number of agents. Normative/prescriptive solutions are sought.

Problem Statement: Let H_0 and H_1 be two alternative hypotheses on the state of the environment and let there be N agents (sensors) who possess some stochastic information related to the state of the environment. In particular, we assume that each agent i observes a random variable y_i with known conditional distribution $P(y_i|H_j)$, j = 0, 1, given either hypothesis. We assume that all agents have information of the same quality, that is, the random variables are identically distributed. Each agent transmits a binary message to a central fusion center, based on his information y_i. The fusion center then takes into account all messages it has received to declare hypothesis H_0 or H_1 true. The problem consists of determining the optimal strategies of the agents as far as their choice of message is concerned. This problem has been long recognized as a prototype problem in team decision theory: it is simple enough so that analysis may be feasible, but also rich enough to allow nontrivial insights into optimal team decision making under uncertainty.

Progress to Date: This problem is being studied by Prof. J. Tsitsiklis. Under the assumption that the random variables y_i are conditionally independent (given either hypothesis), it is known that each agent should choose his message based on a likelihood ratio test. Nevertheless, we have constructed examples which show that even though there is perfect symmetry in the problem, it is optimal to have different agents use different thresholds in their likelihood ratio tests. This is an unfortunate situation, because is severely complicates the numerical solution of the problem (that is, the

explicit computation of the threshold of each agent). Still, we have shown that in the limit, as the number of agents becomes large, it is asymptotically optimal to have each agent use the same threshold. Furthermore, there is a simple effective computational procedure for evaluating this single optimal threshold.

More recently, we showed that if each agent is to transmit K-valued, as opposed to binary messages, then still each agent should use the same decision rule, when the number of agents is large.

We have also investigated the case of M-ary (M > 2) hypothesis testing and obtained evidence indicating that different agents should use different decision rules even in the limit of $N \rightarrow \infty$. The questions concerning the nature of optimal decision rules, as $N \rightarrow \infty$ remains open in this ace.

We also considered a class of decentralized sequential detection problems and showed that only under certain fairly restrictive assumptions do the optimal decision rules have a nice structure.

Documentation: J. Tsitsiklis, "On Optimal Thresholds in Decentralized Detection," submitted for publication to Information and Control; a preliminary version of this paper has been submitted for presentation at the 25th IEEE Conference and Decision and Control, Athens, Greece, December 1986. A formal final draft of this paper will be available in the very near future.

3.3 COMMUNICATION REQUIREMENTS OF DIVISIONALIZED ORGANIZATIONS

Background: In typical organizations, the overall performance cannot be evaluated simply in terms of the performance of each subdivision, as there may be nontrivial coupling effects between distinct subdivisions. These couplings have to be taken explicitly into account; one way of doing so is to assign to the decision maker associated with the operation of each division a cost function which reflects the coupling of his own division with the

remaining divisions. Still, there is some freedom in such a procedure: For any two divisions A and B it may be the responsibility of either decision maker A or decision maker B to ensure that the interaction does not deteriorate the performance of the organization. Of course, the decision maker in charge of those interactions needs to be informed about the actions of the other decision maker. This leads to the following problem. Given a divisionalized organization and an associated organizational cost function, assign cost functions to each division of the organization so that the following two goals are met: a) the costs due to the interaction between different divisions are fully accounted for by the subcosts of each division; b) the communication interface requirements between different divisions are small. In order to assess the communication requirements of a particular assignment of costs to divisions, we take the view that the decision makers may be modeled as boundedly rational individuals, that their decision making process consists of a sequence of adjustments of their decisions in a direction of decreasing costs, while exchanging their tentative decisions with other decision makers who have an interest in those decisions. We then require that there are enough communications so that this iterative process converges to an organizationally optimal set of decisions.

Problem Statement: Consider an organization with N divisions and an associated cost function $J(x_1,\ldots x_N)$, where x_i is the set of decisions taken at the i-th division. Alternatively, x_i may be viewed as the mode of operation of the i-th division. The objective is to have the organization operating at set of decisions $(x_1,\ldots,x_N)$ which are globaly optimal, in the sense that they minimize the organizational cost J. We associate with each division a decision maker DM_i, who is in charge of adjusting the decision unables x_i. We model the decision makers as "boundedly rational" individuals; mathematically, this is translated to the assumption that each decision maker will slowly and iteratively adjust his decisions in a direction which reduces the organizational costs. Furthermore, each decision maker does so based only on partial knowledge of the organizational cost, together with messages received from other decision makers.

Consider a partition $J(x_1,\dots,x_N) = \sum_{i=1}^{N} J^i(x_1,\dots,x_N)$ of the organizational cost. Each subcost J^i reflects the cost incurred to the i-th division and in principle should depend primarily on x_i and only on a few of the remaining x_j's. We then postulate that the decision makers adjust their decisions by means of the following process (algorithm):

(a) DM_i keeps a vector x with his estimates of the current decisions x_k of the other decision makers; also a vector λ with estimates of $\lambda_i^k = \partial J^k/\partial x_i$, for $k \neq i$. (Notice that this partial derivative may be interpreted as DM_i's perception of how his decisions affect the costs incurred to the other divisions.

(b) Once in a while DM_i updates his decision using the rule $x_i := x_i - \gamma \sum_{k=1}^{N} \lambda_i^k$, ($\gamma$ is a small positive scalar) which is just the usual gradient algorithm.

(c) Once in a while DM_i transmits his current decision to other decision makers.

(d) Other decision makers reply to DM_i, by sending a updated value of the partial derivative $\partial J^k/\partial x_i$.

It is not hard to see that for the above procedure to work it is not necessary that all DM's communicate to each other. In particular, if the subcost J^i depends only on x_i, for each i, there would be no need for any communication whatsoever. The required communications are in fact determined by the sparsity structure of the Hessian matrix of the subcost functions J^i. Recall now that all that is given is the original cost function J; we therefore have freedom in choosing the J^i's and we should be able to do this in a way that introduces minimal communication requirements; that is, we want to minimize the number of pairs of decision makers who need to communicate to each other.

The above problem is a prototype organizational design problem and we expect that it will lead to reasonable insights in good organizational structures. On the technical side, it may involve techniques and tools from graph theory. Once the above problem is understood and solved, the next step is to analyze communication requirements quantitatively. In particular, a distributed

gradient algorithm such as the one introduced above converges only if the communication between pairs of DM's should need to communicate are frequent enough. We will then investigate the required frequencies of communications as a function of the strength of coupling between different divisions.

Progress to Date: A graduate student, C. Lee, supervised by Prof. J. Tsitsiklis, has undertaken the task of formulating the problem of finding partitions that minimize the number of pairs of DM's who need to communicate to each other as the topic of his SM research. The literature search phase has been completed, and different problem formulations are being investigated.

Documentation: None as yet.

3.4 COMMUNICATION COMPLEXITY OF DISTRIBUTED CONVEX OPTIMIZATION

Background: The objective of this research effort is to quantify the minimal amount of information that has to be exchanged in an organization, subject to the requirement that a certain goal is accomplished, such as the minimization of an organizational cost function. This problem becomes intereting and relevant under the assumption that no member of the organization "knows" the entire function being minimized, but rather each agent has knowledge of only a piece of the cost function. A normative/prescriptive solution is sought.

Problem Formulation: Let f and g be convex function os n variables. Suppose hat each one of two agents (or decisionmakers) knows the function f (respectively g), in the sense that he is able to compute instantly any quantities associated with this function. The two agents are to exchange a number of binary messages until they are able to determine a point x such that $f(x) + g(x)$ comes within ε of the minimum of $f + g$, where ε is some prespecified accuracy. The objective is to determine the minimum number of such messages that have to be exchanged, as a function of ε and to determine communication protocols which use no more messages than the minimum amount required.

Progress to Date: The problem is being studied by Professor John Tsitsiklis and a graduate student, Zhi-Quan Luo. It is not hard to show that at least $O(n\log 1/\varepsilon)$ messages are needed and a suitable approximate and distributed implementation of ellipsoid-type algorithms demonstrates with $O(n^2\log 1/\varepsilon)$ messages. The challenge is to close this gap. This has been accomplished for the case of one-dimensional problems $n = 1$ for which it has been shown that $O(\log 1/\varepsilon)$ messages are also sufficient. We hope that the technique employed in the one-dimensional case will be generalized for the n-dimensional case, in such way that an algorithm with $O(n^2\log 1/\varepsilon)$ communications will result; we will thus obtain an algorithm which is optimal, as far as the dependence of ε is concerned. The question of the dependence of the amount of communications on the dimension of the problem ($O(n)$ versus $O(n^2)$) seems to be a lot harder and, at present, there are no available techniques for handling it.

Documentation: None as yet, but a presentation on this subject will be made at the 25th IEEE Conference on Decision and Control, Athens, Greece, December 1986.

3.5 DESIGN AND EVALUATION OF ALTERNATIVE ORGANIZATIONAL ARCHITECTURES

Background: The bounded rationality of human decisionmakers and the complexities of the tasks they must perform mandate the formation of organizations. Organizational architectures distribute the decisionmaking workload among the members; different architectures impose different individual loads, lead to different organizational bounded rationality, and result in different organizational performance. Two performance measures have been investigated up to now: accuracy and time delay. An approach to the evaluation and comparison of alternative organizational architectures, that provides insight into the effect structure has on organizational bounded rationality, is the use of a generalized performance-workload locus.

Problem Statement: The development of design guidelines for distributed organizational architectures is the objective. To achieve this objective, a sequence of steps has been defined. Each step in the sequence requires the

solution of both modeling and computational problems:

(1) Development of efficient computational procedures for constructing the generalized performance-workload locus.

(2) Analysis of the functional relationship between internal decision strategies and workload (i.e., the properties of the mapping from strategy space to workload space).

(3) Development of quantitative and qualitative relationships between organizational architecture and the geometry of the performance-workload locus.

Remarks: The work implied in the problem statement requires modeling, analysis, and computation. The use of computer graphics is an integral part of the computational procedures.

Progress To Date: Two subtasks, both addressing organizational delays, were carried out during this quarter. The first subtask consisted of the development and implementation of an algorithm for computing the probability density function of the overall delay when the pdf's of the individual processes are given. The algorithm consists of five steps:

Step 1: Assign pdf's to the transitions representing the situation assessment, information fusion, and response selection processes.

Step 2: Assign pdf's to the delays across communication links.

Step 3: Identify all simple information flow paths and compute the pdf of the delay for each path by convolving the pdf's of the transitions and communications links.

Step 4: Identify all sets of concurrently active paths and obtain the pdf of each set by obtaining the pdf of the maximum delay for each set.

Step 5: Compute the probability of each set of cocurrent paths being active. Then weigh the pdf of each group by the corresponding probability and sum to obtain the pdf of the total delay (or response time).

This algorithm has been implemented in Turbo PASCAL and runs on an IBM PC/AT. In the current implementation, beta pdf's are used to characterize the delays of the transitions and communication links. The algorithm developed by Jin (Thesis #3 in Section 5) is being used to identify the simple paths and to obtain the sets of concurrent paths.

The second subtask addresses the effect that jamming could have on the pdf of the organizational delay. In this case, a very simple model of jamming has been used in order to test the basic concept. Jamming has been described macroscopically by the degree of jamming α that varies between 0 and 1. The effect of jamming has been modeled by modifying the pdf of the delay that characterizes a communication link as follows: (a) the mean of the delay is multiplied by $k = 1(1-\alpha)$ while the variance is multiplied by $\lambda = 1/(1-\alpha)^2$. The effect of jamming is then to increase the average delay and to increase the variance.

These two subtasks lead to the computation of two measures of performance for an organization (a) expected delay, and (b) probability that the response time is less than a threshold value. The threshold could well be the upper bound of the window of opportunity for a given mission. The algorithm for the computation of the two MOPs and the jamming model were used to evaluate the performance of a hierarchical and a parallel organization engaged in an air defense task. It was shown that, in this particular case, the hierarchical and the parallel organizations are comparable for low jamming level. At high jamming levels, however, the hierarchical organization's response time increases substantially.

This work is being carried out by Mr. Stamos Andreadakis under the supervision of Dr. A. H. Levis. In addition, Prof. A. Ephremides of the University of Maryland contributed to the formulation of the jamming problem.

Future Directions: With the basic elements for the computation of measures of performance now in place, the organizational design problem is being formulated by Mr. Andreadakis. This formulation will be based on the

generalized performance-workload locus and will incorporate the analysis and design tools developed thus far in several theses.

Documentation: The results of this work were presented at the 9th MIT/ONR Workshop on C^3 Systems in Monterey, California. The paper is in preparation for inclusion in the Workshop Proceedings. A paper has been accepted for presentation at the 25th IEEE Conference on Decision and COntrol, Athens, Greece, December 1986.

3.6 COMPUTER GRAPHICS FOR ORGANIZATIONAL DESIGN

Background: The analysis of organizations has been based on the ability to construct the performance-workload locus. This is the locus of points that characterize the performance of an organizational form, as described in Section 3.5. This locus serves as the basis for analysis, evaluation, and design of organizations. Indeed, the computer aided design procedure that is being investigated depends on the ability to construct and manipulate this locus with ease.

Problem Statement: There is need to automate the generation of the performance-workload locus. Furthermore, the use of graphics, will allow one to view and compare loci produced from different organizational forms or different values of the design parameters of the same structure. To produce a software system that allows one to do this, several parts needed to be designed, implemented and tested:

(1) A data structure which provides an efficient but general way of storing the data generated by diverse applications.

(2) Software implementation of algorithms for constructing the loci.

(3) Software implementation of algorithms for viewing the loci in different ways, by rotating, translating, or projecting them.

(4) An interactive interface for users to control viewing MOP space.

Remarks: The IBM PC/AT with the Professional Graphics system is being used for this work. This includes the Professional Graphics Controller (PGC), the IBM Professional Graphics Display, the IBM Graphics Toolkit Development Software, and IBM Professional Fortran. The PGC consists of an 8088 microprocessor which executes 3-D graphics routines in ROM (ready-only memory). The graphics routines are executed to perform rotation of the loci about the center point with the amount of rotation specified by the user. This system produces high-resolution color graphics. High quality color plots can be produced on an HP six-pen plotter (Model HP 7475A). A user's manual is now available.

Progress to Date: The development of the software to construct and evaluate the loci was undertaken by Ms. Christine Bohner under the supervision of Dr. A. H. Levis. The development of the first version of the software has been completed. It allows for two-dimensional viewing and three-dimensional viewing with rotation around the x, y, and z axes, when requested. This software is written in Professional Fortran and uses the Graphics Development Toolkit. A simple interface to the program allows the user to control how the loci are viewed. The interface consists of a series of menus and questions asked to the user. For example, when the user requests to change the scale of the loci, the program lists the present scale factors and prompts for new factors in the x, y, and z axes.

The data structure is key in defining how the data are to be displayed. It has the following construction:

$$LOCI[I,J,K,L \,||\, x_1,x_2,x_3,x_4,x_5,x_6,x_7]$$

$x_1,\ldots x_7$ represent the actual data points in the loci. Each x represents one dimension of the locus space. The user can choose which two or three dimensions are to be viewed on the screen and thus obtain different projections of the loci. When the software is applied to

organizational design, the x's represent the accuracy measure, the time delay measure, and the workload of each decisionmaker. The current system can handle organizations with up to five members (n=5).

I,J,K,L are indices that represent the variables that are varied in a specified order to generate the loci. The user can select the order in which these will vary and thus change how the loci are drawn.

The original version developed by Ms. Bohner, which is a general purpose tool, has been enhanced by S. Andreadakis for application to the organizational design problem. The software has been used already for the analysis of the three-person hierarchical and parallel organizations mentioned in Section 3.5.

Future Work: As the software is being used, enhancements become necessary. The first enhancement will be the creation of a set of axes that are properly scaled and that follows the commands issued for transforming the locus (rotate, translate, project, etc.).

A second enhancement will address the user interface. We expect to use the Screen Sculptor (TM) to construct an interface that allows the user to control the presentation of the locus.

Documentation: The basic design has been described in Ms. C. Bohner's thesis, "Computer Graphics for System Effectiveness Analysis," MS Thesis, LIDS-TH-1573, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, June 1986.

3.7 ASYNCHRONOUS PROTOCOLS AND ORGANIZATIONAL STRUCTURES

Background: A key driving force in the design of distributed decisionmaking architectures is the rate of change of the environment, i.e., the tempo of operations. As the tempo of operations has been increasing, the time available from the moment an event takes place until effective action has

been taken and completed has been decreasing. The ability to respond reliably to events in a timely manner, i.e., before responses are pre-empted, is characterized by the performance measure of "timeliness". Analytical techniques are needed for evaluating this performance measure and for designing organizational architectures and associated protocols that optimize timeliness. Several research problems have been defined and are being pursued. The one on which most progress was achieved during the past quarter is:

Organizational Structures

Problem Statement: The objective of this task is the development of analytical and computational tools for the design of organizational structures. The emphasis of this task is in the generation of alternative topologies that satisfy task requirements.

Progress to Date: This task is being carried out by P. Remy and V. Jin under the supervision of Dr. A. H. Levis. The first step in the procedure was the definition of the Petri Net and the corresponding data structure for the interacting decisionmaker. In the past, information sharing was allowed only between the situation assessment stage and the information fusion process. This assumption has been relaxed to allow four different forms of information sharing - each form depends on the source of the information (e.g., is one DM informing the other of his situation assessment or of his response?) and on the destination. For example, the situation assessment of one DM may be the input to the next one in a serial or hierarchical organization.

After defining the set of possible interactions, a combinatorial problem is formulated. The dimensionality of this problem is prohibitive, if no constraints on the structure are imposed. (There are $2^{4n(n-2)}$ organizational forms in this formulation, where n is the number of decisionmakers) However, an algorithmic approach has been developed by P. Remy that reduces the problem to a computationally tractable one.

This approach is based on several notions: (a) the use of constraints that are specific to the application, (b) the use of generic structural constraints, (c) the concept of a simple path (or S-invariants of a Petri Net), and (d) the concept of concurrent paths.

The preliminary results of this investigation have been documented in a paper by P. Remy, A. H. Levis, and V. Jin (Appendix I).

Future Work: The algorithm for generating feasible organizational structures will be implemented to accommodate any five person organization. The algorithm developed by Jin will be used to check the violation of constraints; an alternative algorithm based on the algebraic properties of S-invariants developed by Martinez and Silva as modified by Toudic will also be implemented in order to check results and compare the efficiency of the two approaches.

Documentation: P. Remy, A. H. Levis, V. Jin, "On the Design of Distributed Organizational Structures," LIDS-P-1581, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, July 1986.

This paper was presented at the 9th MIT/ONR Workshop on C^3 Systems. It has also been submitted for presentation at the X IFAC World Congress in Munich, FRG in 1987.

4. RESEARCH PERSONNEL

Prof. Michael Athans,	Co-principal investigator
Dr. Alexander H. Levis,	Co-principal investigator
Prof. John Tsitsiklis	
Mr. Stamatios Andreadakis,	graduate research assistant (Ph.D.)
*Ms. Chris Bohner,	graduate research assistant (M.S)
Mr. Herve Hillion	graduate research assistant (M.S)
Ms. Victoria Jin,	graduate research assistant (Ph.D.)
Ms. Zhi Quan Luo	graduate research assistant (M.S)
*Mr. Jason Papastavrou,	graduate research assistant (M.S)
Mr. Pacasl Remy	graduate research assistant (M.S)

*Received M.S. degree.

5. DOCUMENTATION

5.1 Theses

[1] J. N. Tsitsiklis, "Problems in Decentralized Decision Making and Computation," Ph.D. Thesis, Dept. of EECS also Report LIDS-TH-1424, Laboratory for Information and Decision Systems, MIT, Cambridge, Mass., December 1984.

[2] G. J. Bejjani, "Information Storage and Access in Decisionmaking Organizations," MS Thesis, Technology and Policy Program, also Report LIDS-TH-1434, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, January 1985.

[3] V. Y.-Y. Jin, "Delays for Distributed Decisionmaking Organizations," MS Thesis, Mechanical Engineering, also Report LIDS-TH-1459, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, May 1985.

[4] K. L. Boettcher, "A Methodology for the Analysis and Design of Human Information Processing Organizations," Ph.D. Thesis, Dept. of Electrical Engineering and Computer Sciences, also Report LIDS-TH-1501, Laboratory for Information and Decision Systems, MIT, September 1985.

[5] R. P. Wiley, "Performance Analysis of Stochastic Timed Petri Nets," Ph.D. Thesis, Dept. of Electrical Engineering and Computer Sciences, also, Report LIDS-TH-1525, Laboratory for Information and Decision Systems, MIT, January 1986.

[6] J. D. Papastavrou, "Distributed Detection with Selective Communications," MS Thesis, Dept. of Electrical Engineering and Computer Sciences, also Report LIDS-TH-1563, Laboratory for Information and Decision Systems, MIT, May 1986.

[7] C. M. Bohner, "Computer Graphics for Systems Effectiveness Analysis," MS Thesis, Technology and Policy Program, also Report LIDS-TH-1573, Laboratory for Information and Decision Systems, MIT, June 1986.

5.2 Technical Papers

[1] R. P. Wiley and R. R. Tenney, "Performance Evaluation of Stochastic Timed Decision-Free Petri Nets," Proc. 24th IEEE Conference on Decision and Control, December 11-13, 1985, Ft. Lauderdale, Florida, pp. 558-563.

[2] G. J. Bejjani and A. H. Levis, "Information Storage and Access in Decisionmaking Organizations," LIDS-P-1466, Laboratory for Information and Decision Systems, MIT, May 1985.

[3] V. Y.-Y. Jin and A. H. Levis, "Computation of Delays in Acyclical Distributed Decisionmaking Organizations," LIDS-P-1488, Laboratory for Information and Decision Systems, MIT, August 1985.

[4] K. L. Boettcher and R. R. Tenney, "A Case Study in Human Team Decisionmaking, Proc. Systems, Man, and Cybernetics Symposium, Tucson, AZ, November 1985.

[5] K. L. Boettcher and R. R. Tenney, "On the Analysis and Design of Human Information Processing Organizations," LIDS-P-1503, Proc. 8th MIT/ONR Workshop on C^3 Systems, December 1985.

[6] K. L. Boettcher and R. R. Tenney, "Distributed Decisionmaking with Constrained Decision Makers - A Case Study," Proc. 8th MIT/ONR Workshop on C^3 Systems, December 1985.

[7] R. P. Wiley and R. R. Tenney, "Calculating Time-Related Performance Measures of a Distributed Tactical Decisionmaking Organization Using Stochastic Timed Petri Nets," Proc. 8th MIT/ONR Workshop on C^3 Systems, December 1985.

[8] V. Y.-Y. Jin, A. H. Levis, and P. A. Remy, "Delays in Acyclical Distributed Decisionmaking Organizations," LIDS-P-1528, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, also in Proc. 4th IFAC/IFORS Symposium on Large Scale Systems: Theory and Applications, Zurich, Switzerland, August 1986.

[9] J. Papastavrou and M. Athans, "A Distributed Hypotheses-Testing Team Decision Problem with Communications Cost," LIDS-P-1538, Laboratory for Information and Decision Systems, MIT, February 1986, to appear in Proc. of 25th IEEE Conference on Decision and Control, December 1986.

[10] J. Tsitsiklis, "On Optimal Thresholds in Decentralized Detection," to appear in the Proc. of 25th IEEE Conference on Decision and Control, December 1986.

[11] P. Remy, A. H. Levis, V. Jin, "On the Design of Distributed Organizational Structures," LIDS-P-1581, Laboratory for Information and Decision Systems, MIT, July 1986, to appear in Proc. of 25th IEEE Conference on Decision and Control, December 1986.

ON THE DESIGN OF DISTRIBUTED ORGANIZATIONAL STRUCTURES*

by

Pascal A. Remy

Alexander H. Levis

Victoria Y.-Y. Jin

ABSTRACT

The problem of designing human decisionmaking organizations is formulated as an organizational form problem with special structure. Petri Nets are used to represent the organizational form. An algorithmic procedure, suitable for computer-aided design, is presented and the specific algorithms that it includes are developed. The approach reduces the dimensionality of the problem to a tractable level.

*This work was supported by the U.S. Office of Naval Research under Contract No. N00014-84-K-0519 and by the Joint Directors of Laboratories through the Office of Naval Research Contract No. N00014-85-K-0782.

ON THE DESIGN OF DISTRIBUTED ORGANIZATIONAL STRUCTURES+

Pascal A. Remy, Alexander H. Levis*, and Victoria Y.-Y. Jin

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Room 35-410
Cambridge, MA 02139, USA

Telex: 92-1473
Telephone: (617) 253-7262

Subject Area: 12.1 Large Scale Systems

Keywords: Organizations, Distributed Decision Making,
Hierarchical Structures, Parallel Structures

ABSTRACT

The problem of designing human decisionmaking organizations is formulated as an organizational form problem with special structure. Petri Nets are used to represent the organizational form. An algorithmic procedure, suitable for computer-aided design, is presented and the specific algorithms that it includes are developed. The approach reduces the dimensionality of the problem to a tractable level.

---

*Corresponding Author
+This work was supported by the U.S. Office of Naval Research under Contract No. N00014-84-K-0519 and by the Joint Directors of Laboratories Contract No. ONR/N00014-85-K-0782.

INTRODUCTION

There are two basic problems in organizational design: the problem of organizational form and the problem of organizational control. Most of the theoretical developments in decision and control theory have addressed the latter problem: given an organizational structure, determine the decision rules or strategies that optimize some performance criterion. The former problem has been addressed only indirectly, i.e., given an organizational form, evaluate its performance according to some criteria and then change in some ad hoc manner the organizational form until a satisfactory structure has been obtained. The reason for this approach is that the general organizational form problem becomes computationally infeasible, even for a small number of organizational units.

In this paper, the organizational form problem is posed for a well defined class of organizations - those that have fixed structure and can be represented by acyclical directed graphs. These structures represent distributed decisionmaking organizations performing well defined tasks under specified rules of operation. Such organizations have been modeled and analyzed in a series of papers [1-4]. The basic unit of the models is the interacting decisionmaker with bounded rationality. The set of interactions will be generalized in Section 2 to allow not only for information sharing and command inputs, but also several forms of result sharing between decisionmakers. While this generalization increases the dimensionality of the design problem, it also allows for more realistic models of actual organizational interactions.

The mathematical formulation of the problem is based on the Petri Net description of the organizational structure. Furthermore, the dimensionality of the combinatorial problem is reduced by utilizing the notion of information paths within the organization. A number of new concepts are introduced that bound the problem to the search for alternative organizational forms from within the set of feasible structures only. The introduction of structural constraints, which characterize the

class of organizations under consideration, and of user constraints that are application specific lead to an algorithmic approach that is implementable on a personal computer. The mathematical model of the organization is described in the second section. In the third section, the various constraints are introduced. In the fourth section, the algorithm is discribed, while in the fifth a nontrivial example is presented.

MATHEMATICAL MODEL

The single interacting decisionmaker is modeled as having four stages or actions, the situation assessment (SA) stage, the information fusion (IF) stage, the command interpretation (CI) stage, and the response selection (RS) stage. In the SA stage, external inputs -- data from the environment or other members of the organization are processed to determine the situation assessment. This information is transmitted to the IF stage where it is fused with situation assessments communicated by other organization members. The resulting revised situation assessment is used to select response in the response selection stage. The responses can be restricted by commands received by the CI stage that precedes the RS stage. An individual decisionmaker could receive inputs therefore at the SA stage, the IF stage, and the CI stage. It can produce outputs only by the SA stage and the RS stage. The exchange of information between the situation assessment and the information fusion stages of different decisionmakers constitute <u>information sharing</u> among them. On the other hand, what is being transmitted from the response selection stage of one decisionmaker (DM) to the IF stage of another could be the decision made by the first DM; in this case, the interaction is of the <u>result sharing</u> type. If the transmission is from the RS stage of one to the CI stage of another, then the former is issuing a command to the latter. This interaction imposes a hierarchical relationship between decisionmakers, - one is a commander, the other is a subordinate - while the other interactions don't.

The use of Petri Nets for the modeling of decisionmaking organizations was presented in [3] and exploited in [4]. Petri Nets [5] are bipartite

directed multigraphs. The two types of nodes are places, denoted by circles and representing signals or conditions, and transitions, denoted by bars and representing processes or events. Places can be connected by links only to transitions, and transitions can be connected only to places. The links are directed. Tokens are used to indicate when conditions are met - tokens are shown in the corresponding place nodes. When all the input places to a transition contain tokens, then the transition is said to be enabled and it can then fire. Properties of Petri Nets are the subject of current research, e.g., references [5] - [8].

Figure 1 shows the Petri Net model of the single interacting decisionmaker. The DM can receive inputs (u) only at the SA, IF, and CI stages and produce outputs (y) only by the SA and RS stages, as stated earlier.

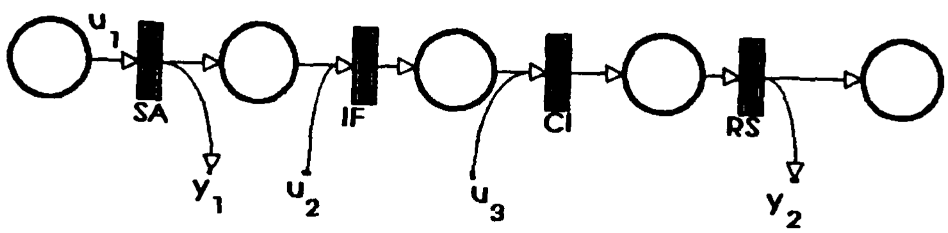

Figure 1. Aggregated Model of Interacting Decisionmaker

The allowable interactions between two decisionmakers are shown in Figure 2. For clarity, only the interactions from DM^i to DM^j are shown. The interactions from DM^j to DM^i are identical. The superscripts i or j denote the decisionmaker; the pair of superscripts ij indicates a link from DM^i to DM^j. Consider the general case of an organization consisting of N decisionmakers, a single input place, and a single output place. The last two are not really restrictions, since multiple sources can be represented by one place and a transition that partitions the input and distributes it to the input places of the appropriate organization members.

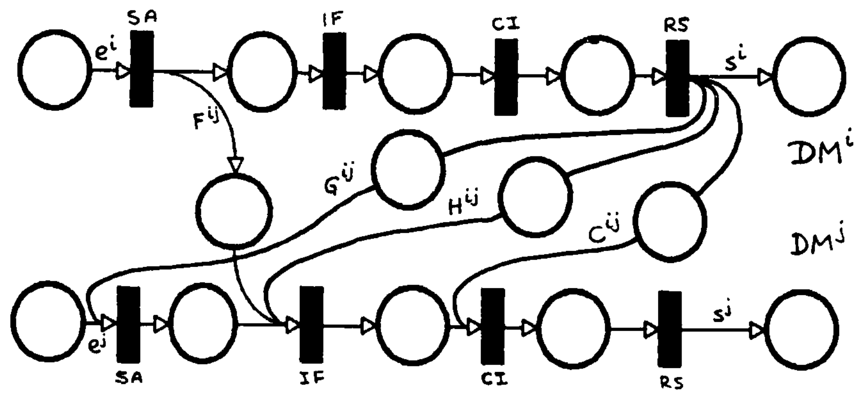

Figure 2. Modeled Interactions Between Two Decisionmakers

The organizational structure, as depicted by the Petri Net, can be expressed in terms of two vectors and four matrices. The elements of these vectors and matrices can take the value of zero or of one; if zero, then there is no connection, if one, then there is.

The interaction between the organization and the external source (input) is represented by an N-dimensional vector $\underline{e}$ with elements e^i. The output from the RS stage to the external environment is represented by the N-dimensional vector $\underline{s}$ with elements s^i.

The information flow from the SA stage of DM^i to the IF stage of DM^j is denoted by F^{ij}. Since each DM can share situation assessment information with the other N-1 DMs, the matrix F is N x N, but with the diagonal elements identically equal to zero.

Similarly, the links between the RS stage of a DM and the SA stage of the others are represented by the matrix G; the links from the RS stage to the IF stage by H; and the links from the RS stage to the CI stage by the matrix C. These three matrices are also N x N and their diagonal elements are identially equal to zero.

Therefore

$$\underline{e} = [e^i],\ \underline{s} = [s^i] \qquad 1 \le i \le N\ ,\ 1 \le j \le N$$

$$F = [F^{ij}],\ G = [G^{ij}],\ H = [H^{ij}],\ C = [C^{ij}]$$

$$F^{ii} = G^{ii} = H^{ii} = C^{ii} = 0,\ \text{all } i$$

There are, altogether, 2^m possible combinations of different vectors $\underline{e}$, $\underline{s}$ and matrices $\underline{F}$, $\underline{G}$, $\underline{H}$, and $\underline{C}$, where $m = 4N^2 - 2N$. For a five member organization ($N=5$), m is equal to 90 and the number of alternatives is 2^{90}. Fortunately, many of these are not valid organizational forms and need not be considered. In the next section, the allowable combinations will be restricted by defining a set of structural constraints.

CONSTRAINTS

Four different structural constraints are formulated that apply to all organizational forms being considered.

R_1 The structure should have no loops.

R_{2a} The structure should be connected, i.e., there should be at least one undirected path between any two nodes in the structure.

R_{2b} A directed path should exist from the source to every node of the structure and a directed path should exist from any node to the output node.

R_3 There can be at most one link from the RS stage of a DM to each one of the other DMs, i.e., for each i and j, only one of the triplet (G^{ij}, H^{ij}, C^{ij}) can be nonzero.

R_4 Information fusion can take place only at the IF and CI stages, consequently, the SA stage of each DM can have only one input from outside of the DM.

The set of structural constraints is defined as

$$R_s = \{R_1,\ R_{2a},\ R_{2b},\ R_3,\ R_4\}$$

The first constraint allows acyclical organizations only. The second and third define connectivity as it pertains to this problem; it eliminates structures that do not represent a single organization. The last two reflect the meaning of the four-stage decisionmaking model.

In addition to these constraints, the organization designer may introduce additional ones that reflect the specific application he is considering. For example, there may be a hierarchical relationship between the decisionmakers that must be maintained in the organizational structure. Then, the appropriate 0s and 1s will be placed in the arrays {e,s,F,G,H,C,} thus restricting even further the organizational design problem solution. Let these constraints be denoted by R_u.

A Petri Net whose structure can be modeled by the four matrices and two vectors {F,G,H,C} and {e,s}, respectively will be called a Well Defined Net (WDN). A WDN that fulfills the structural constraints R_s will be called an <u>Admissible</u> organizational form, while WDNs that satisfy both the structural <u>and</u> the designer's contraints will be called a <u>Feasible</u> organization form. A decision tree showing the relationship between the different sets of organizational forms is shown in Figure 3.

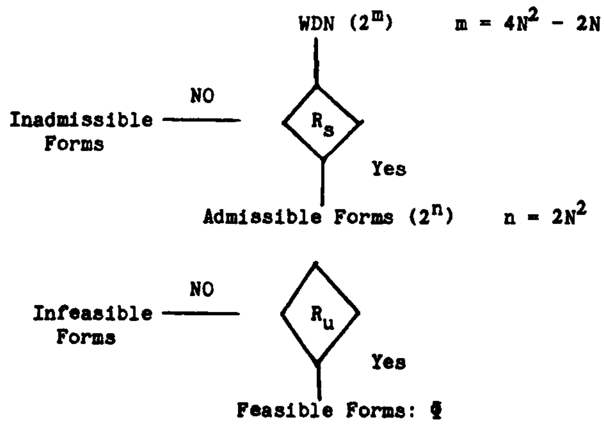

Figure 3. Organizational Forms

The notion of a subnet of well defined nets (WDNs) can be defined as follows: Let W be a WDN specified by the set of arrays $\{\underline{e},\underline{s},F,G,H,C\}$. Let W′ be a second WDN specified by the set $\{\underline{e}',\underline{s}',F',G',H',C'\}$. Then W′ is a subnet of W if and only if

$$\underline{e}' \leq e \quad , \quad \underline{s}' \leq \underline{s} \quad , \quad D' \leq F$$

$$G' \leq G \quad , \quad H' \leq H \quad , \quad C' \leq C$$

where the inequality between arrays means that

$$(A' \leq A) \qquad (\forall i \ , \ \forall j \quad A'_{ij} \leq A_{ij}).$$

Therefore, W′ is a subnet of W if any interaction in W′ (i.e., a 1 in any of the arrays e′,s′,F′,G′,H′,C′) is also an interaction in W (i.e., a 1 in the corresponding array of W). The union of two subnets W_1 and W_2 is a new net that contains all the interactions that appear in either W_1 or W_2 or both.

DESIGN ALGORITHM

Let R be the set of contraints $R_s \cup R_u$. The design problem is to determine all the Feasible Organizational Forms, $\Phi(R)$, i.e., all the WDNs that satisfy the set of contraints R. The approach is based on defining and constructing two subsets of feasible organizational forms: the maximally connected organizations and the minimally connected organizations.

A Feasible Organizational form is a Maximally Connected Organization (MAXO) if and only if it is not possible to add a single link without violating the constraint set R. The set of MAXOs will be denoted by $\Phi_{max}(R)$.

A Feasible Organizational form is a Minimally Connected Organization

(MINO) if and only is it is not possible to remove a single link without violating the constraint set R. The set of MINOs is denoted by $\Phi_{min}(R)$.

Consider now the designer's constraints R_u. The well defined nets that satisfy the constraints R_u are denoted by the set $\Omega(R_u)$. For a given number of decisionmakers, the maximally connected net associated with the set of constraints R_u is obtained by replacing all the undetermined elements of $\{\underline{e},\underline{s},F,G,H,C\}$ with 1s. This particular net is denoted by $\tilde{\Omega}(R_u)$. Therefore, by construction, $\tilde{\Omega}(R_u)$ is unique.

<u>Proposition 1</u>: Any feasible organization $\Phi(R)$ is a subnet of $\tilde{\Omega}(R_u)$.

Since any element of $\Phi(R)$ must satisfy the set of constraints R_u and since $\tilde{\Omega}(R_u)$ is the MAXO with respect to R_u, the elements of $\Phi(R)$ must be subnets of $\tilde{\Omega}(R_u)$.

Since $\tilde{\Omega}(R_u)$ is a Petri Net, it has an associated incidence (or flow) matrix Δ[4]. The rows of the incidence matrix represent the places, while the columns represent the transitions. A -1 in position Δ_{ij} indicates that there is a directed link from place i to transition j; a +1 indicates a directed link from transition j to place i, while a 0 indicates the absence of a directed link between place i and transition j.

An integer vector $\underline{q}$ is an s-invariant of $\tilde{\Omega}(R_u)$ if and only if

$$\Delta'\underline{q} = 0$$

A <u>simple</u> information <u>path</u> of $\tilde{\Omega}(R_u)$ is a minimal support s-invariant of $\tilde{\Omega}(R_u)$ that includes the source node (source place) (for details, see [4]). This simple path is a directed path without loops from the source of the net to the sink.

<u>Proposition 2</u>: Any well defined net that satisfies the constraints R_u and the connectivity constraint R_{2b} is a union of simple paths of $\Omega_{max}(R_u)$.

Proof: If a WDN Υ satisfies the constraint set R_u, then it is a subnet of $\tilde{\Omega}(R_u)$, by the definition of $\tilde{\Omega}(R_u)$. Constraint R_{2b} implies that every node of Υ is included in at least one simple path since there is a path from the source to the node and a path from the node to the output node. Therefore, Υ is a union of simple paths of $\tilde{\Omega}(R_u)$.

Corollary: Any feasible organization Φ is a union of simple paths of $\tilde{\Omega}(R_u)$.

Let $Sp(R_u)$ be the set of all simple paths of $\tilde{\Omega}(R_u)$, i.e.,

$$Sp(R_u) = \{sp_1, sp_2, \ldots, sp_r\}$$

and let $\cup Sp(R_u)$ denote the set of all unions of simple paths of $\tilde{\Omega}(R_u)$. If there are r elements in $Sp(R_u)$ then there are 2^r elements in $\cup Su(R_u)$. From now on, only WDNs that are elements of $\cup Sp$ need be considered.

The procedure described so far can be summarized by a sequence of four steps.

Step 1: Given the set of contraints R_u, define the set of arrays {e,s,F,G,H,C} that satisfy these constraints.

Step 2: Construct the maximally connected net $\tilde{\Omega}(R_u)$ by replacing with 1s all the undetermined elements in the six arrays.

Step 3: Find all the simple paths of $\tilde{\Omega}(R_u)$ using the algorithms described in [3] or the algorithm of Martinez and Silva [9] which generates all minimal support s-invariants of a general Petri Net using linear algebra tools. An improved version of this algorithm has been proposed by Toudic [10].

Step 4: Construct the set of all unions of simple paths of $\tilde{\Omega}(R_u)$.

From the corollary, the set $\{\Phi\}$ is a subset of $\cup Sp(R_u)$. Consequently, the number of feasible organizational forms is bounded by 2^r. The dimensionality of the problem is still too large. Once more step is needed to reduce the computational effort.

<u>Proposition 3</u>: Let $\mathbf{T}$ be a WDN that is a union of simple paths of $\tilde{\Omega}(R_u)$. Then $\mathbf{T}$ is a feasible organization form, i.e., $\mathbf{T} \in \{\Phi\}$, if and only if, (a) there is at least one MINO which is a subnet of $\mathbf{T}$, and (b) $\mathbf{T}$ is the subnet of at least one MAXO.

The MAXOs and MINOs can be thought of a the "boundaries" of the set $\{\Phi\}$. The next step is to find a procedure for constructing the MAXOs and the MINOs corresponding to the constraint set R. Since $\mathbf{T}$ is a subset of $\cup Sp(R_u)$, it follows that Φ_{min} is a subset of $\cup Sp(R_u)$. Then, one can scan all the elements of $\cup Sp$ and select those that satisfy the constraints R. A tree-structure is defined to guide the search, Figure 4.

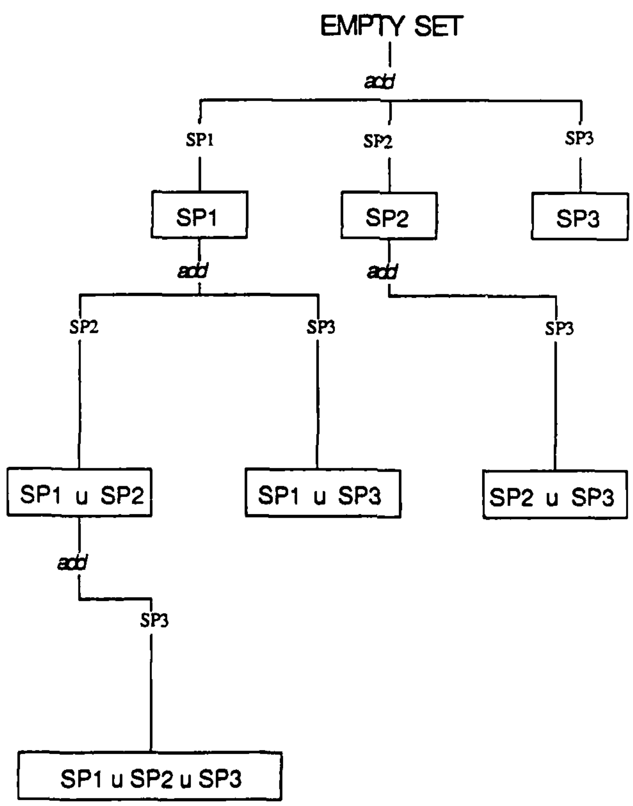

The tree is scanned from Right to Left and from the top down. For the case r=3 (Fig. 4), the search for the MINOs starts as follows:

$$sp_3 \rightarrow sp_2 \rightarrow sp_1 \qquad\qquad sp_3 \cup sp_2 \rightarrow sp_3 \cup sp_1 \rightarrow \ldots$$

At each step, the constraints R are checked to see if they are violated. If they are, the algorithm goes to the next subnet. If the structural constraints are satisfied at some point in the search, then the subnet is a MINO. At that point, all branches of the tree that contain the MINO just found are eliminated. The algorithm proceeds by continuing the search on the "pruned" tree until it finds another MINO. The procedure is repeated until there are no more subnets to be checked.

To determine the MAXOs, the same procedure is used but instead of creating a tree starting with the null subnet and building the subnets by taking the union of simple paths, the tree has as its root the net $\tilde{\Omega}(R_u)$. The tree is constructed by removing, in sequence, one path, two paths, etc. The tree is scanned again from right to left, and from the top down. When a MAXO is found, it is kept and the corresponding branches are eliminated.

Therefore, the fifth and sixth steps of the algorithm are:

Step 5: Construct the tree of the set $\cup$Sp with the null set at the root and search to find the minimally connected organizations.

Step 6: Construct the tree of the set $\cup$Sp with the $\tilde{\Omega}(R_u)$ at the root and search to find the maximally connected organizations.

Implicit in Steps 5 and 6 is the ability to test efficiently whether constraints R are satisfied. Indeed, if the interconnection matrix (see Ref. [4]) for the net $\tilde{\Omega}(R_u)$ is constructed, then the checking for the constraints R reduces to simple tests on the elements of the interconnection matrix.

In this paper, the procedure is illustrated for the case of a two person organization; the application of interest is a five person organization. The design of the five person organization constitutes work in progress.

APPLICATION

Consider a two person organization (N=2). In general, there are 2^{13} well defined nets for this problem. Clearly, even for this very small example a direct approach would not be practical!

A computer-aided design procedure has been implemented on an IBM PC/AT with 512K RAM and a 20 MB hard disk drive. The six arrays for organizations with up to 5 members are shown graphically on the color monitor; a simplified printout of the screen can be obtained (Fig. 5). The symbol # denotes that no link can exist at this location. A 0 indicates the choice that no link be at that location, a 1 that a link will be at that location, and an x indicates that the choice is open: the x's represent the degrees of freedom in the design.

```
****************************************************************************
*                   ORGANIZATIONAL   FORM   DESIGN   -   General case
****************************************************************************
*                      1  2  3  4  5                           1  2  3  4  5      *
*                    .................                       .................    *
*  e: input      0   . 1  1  1  0  0 .      s: output    0   . 0  0  1  1  1 .    *
*                    .................                       .................    *
*                      1  2  3  4  5                           1  2  3  4  5      *
*                    .................                       .................    *
*     F          1   . #  x  x  x  x .         G         1   . #  0  0  x  x .    *
*                2   . x  #  x  x  x .                   2   . 0  #  0  x  x .    *
*  SA - IF       3   . x  x  #  x  x .      RS - SA      3   . 0  0  #  x  x .    *
*                4   . x  x  x  #  x .                   4   . 0  0  0  #  x .    *
*                5   . x  x  x  x  # .                   5   . 0  0  0  x  # .    *
*                    .................                       .................    *
*                      1  2  3  4  5                           1  2  3  4  5      *
*                    .................                       .................    *
*     H          1   . #  0  0  0  0 .         C         1   . #  x  x  1  1 .    *
*                2   . 0  #  0  0  0 .                   2   . 0  #  0  1  1 .    *
*  RS - IF       3   . 0  0  #  0  0 .      RS - CI      3   . 0  x  #  x  x .    *
*                4   . 0  0  0  #  0 .                   4   . 0  0  0  #  0 .    *
*                5   . 0  0  0  0  # .                   5   . 0  0  0  0  # .    *
*                    .................                       .................    *
****************************************************************************
Press [PgUp],[PgDn]=Prior,Next Screen; [Esc]=Exit; or Any Other=Continue:
```

Figure 5. Simplified Representation of the Screen

The designer introduces the 0's and 1's that represent the constraint set R_u. For example, the designer has decided that both decisionmakers will receive inputs from the environment, that only the second one will produce an output, and that the two cannot issue commands to each other. The resulting six arrays that incorporate the constraints R_u are (Step 1).

$$\underline{e} = [\ 1\ 1\ ] \qquad \underline{s} = [\ 0\ 1\ ]$$

$$F = \begin{bmatrix} \# & x \\ x & \# \end{bmatrix} \quad G = \begin{bmatrix} \# & x \\ x & \# \end{bmatrix} \quad H = \begin{bmatrix} \# & x \\ x & \# \end{bmatrix} \quad C = \begin{bmatrix} \# & 0 \\ 0 & \# \end{bmatrix}$$

The first constraint on the input resulted in specifying the array e; the constraint on the output the array $\underline{s}$. The requirements that neither can issue commands resulted in array C being completely specified. The remaining degrees of freedom for this design are six so that the total number of possible nets in Ω is 2^6. The maximally connected net $\tilde{\Omega}$ is obtained by setting all x's equal to 1 (Step 2). The next step (#3) requires the determination of all the simple paths in $\bar{\Omega}$. There are three (r=3) simple paths; they are shown in Fig. 6. Given that r=3, the tree for determining the MINOs is the one shown in Fig. 4 which contains all the unions of the three paths. Carrying out the search of this tree leads to three MINOs, namely, sp_1, sp_2, and sp_3, i.e., each simple path is a feasible organization, an element of the set $\{\Phi\}$ (Step 5). Finally, Step 6, the tree determining the MAXOs is constructed. The search leads to the conclusion that the only two MAXOs are (sp_1 sp_3) and (sp_2 sp_3).

In this very simple case, since there are no intermediate levels in the trees, the complete set $\{\Phi\}$ has been determined: it consists of the three MINOs and the two MAXOs.

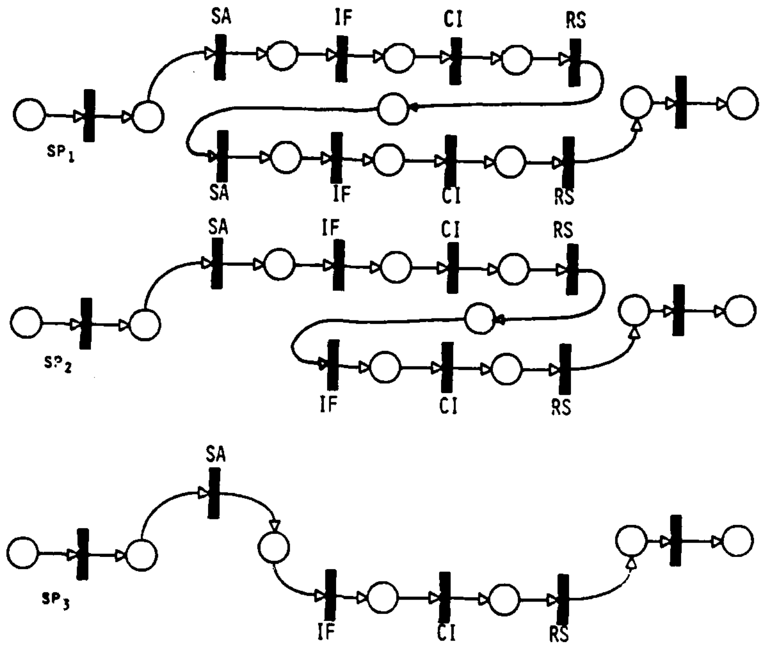

Figure 6. The Three Simple Paths

CONCLUSION

The organizational from problem has been described and a mathematical formulation based on Petri Nets has been presented. An algorithm that reduces the problem to a tractable level has been introduced that takes into account the special structure of human decisionmaking organizations. A preliminary implementation of the algorithm on a microcomputer is described.

REFERENCES

[1] K. L. Boettcher and A. H. Levis, "Modeling and Analysis of Teams of Interacting Decisionmakers with Bounded Rationality," Automatica, Vol. 19, No. 5, November 1983.

[2] A. H. Levis, "Information Processing and Decisionmaking Organizations: A Mathematical Description," Large Scale Systems, Vol. 7, 1984.

[3] D. Tabak and A. H. Levis, "Petri Net Representation of Decisions Models," IEEE Trans. on Systems, Man, and Cybernetics, Vol. SMC-15, No. 6, November/December 1985.

[4] V. Y.-Y. Jin, A. H. Levis, and P. Remy, "Delays in Acyclical Distributed Decisionmaking Organizations," Proc. 4th IFAC/IFORS Symposium on Large Scale Systems, Zurich, Switzerland, August 1986.

[5] J. L. Peterson, Petri Net Theory and the Modeling of Systems, Prentice-Hall, Englewood Cliffs, NJ, 1981.

[6] G. W. Brams, Reseaux de Petri: Theorie et Pratique, Masson, Paris, 1983.

[7] W. Brauer, Net Theory and Applications, Proceedings of the Advanced Course on General Net Theory of Processes and Systems, Hamburg 1979, Springer-Verlag #84, Berlin, 1980.

[8] G. Rozenberg, Advances in Petri Nets 1984, Springer-Verlag, Berlin, 1985.

[9] J. Martinez and M. Silva, A Simple and Fast Algorithm to Obtain all Invariants of a Genralized Petri Net, Giraud C. and Reisig W., Springer-Verlag #52, Berlin, 1980.

[10] J. M. Toudic, Algorithmes d'analyse Structurelle des Reseaux De Petri, These 3 Cycle, Universite Paris VI, 1981.

END

DATE
FILMED

9-86